This Page is Inserted by IFW Indexing and Scanning Operations and is not part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.
As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

IEEE HOME | SEARCH IEEE | SHOP | WEB ACCOUNT | CONTACT IEEE

IEEE

Membership Publications/Services Standards Conferences Careers/Jobs

IEEE Xplore® RELEASE 1.8

Welcome
United States Patent and Trademark Office

Help FAQ Terms IEEE Peer Review Quick Links

» Se

Welcome to IEEE Xplore®
- Home
- What Can I Access?
- Log-out

Tables of Contents
- Journals & Magazines
- Conference Proceedings
- Standards

Search
- By Author
- Basic
- Advanced

Member Services
- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

IEEE Enterprise
- Access the IEEE Enterprise File Cabinet

Your search matched **0** of **1058483** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or enter new one in the text box.

memory hierarchy level<and>library program Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

Results:
No documents matched your query.

Print Format

Home | Log-out | Journals | Conference Proceedings | Standards | Search by Author | Basic Search | Advanced Search | Join IEEE | Web Account | New this week | OPAC Linking Information | Your Feedback | Technical Support | Email Alerting | No Robots Please | Release Notes | IEEE Online Publications | Help | FAQ | Terms | Back to Top

Copyright © 2004 IEEE — All rights reserved

IEEE HOME | SEARCH IEEE | SHOP | WEB ACCOUNT | CONTACT IEEE

IEEE

Membership Publications/Services Standards Conferences Careers/Jobs

IEEE Xplore® RELEASE 1.8

Welcome
United States Patent and Trademark Office

Help FAQ Terms IEEE Peer Review Quick Links

» Se

Welcome to IEEE Xplore®
- Home
- What Can I Access?
- Log-out

Tables of Contents
- Journals & Magazines
- Conference Proceedings
- Standards

Search
- By Author
- Basic
- Advanced

Member Services
- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

IEEE Enterprise
- Access the IEEE Enterprise File Cabinet

Print Format

Your search matched **4** of **1058483** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or enter new one in the text box.

memory hierarchy levels Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

1 **The impact of recent DRAM architectures on embedded systems performance**
Gries, M.;
Euromicro Conference, 2000. Proceedings of the 26th , Volume: 1 , 5-7 Sept.

Pages:282 - 289 vol.1

[Abstract] [PDF Full-Text (696 KB)] **IEEE CNF**

2 **Eager writeback-a technique for improving bandwidth utilization**
Lee, H.H.S.; Tyson, G.S.; Farrens, M.K.;
Microarchitecture, 2000. MICRO-33. Proceedings. 33rd Annual IEEE/ACM International Symposium on , 10-13 Dec. 2000
Pages:11 - 21

[Abstract] [PDF Full-Text (956 KB)] **IEEE CNF**

3 **Optimizing software data prefetches with rotating registers**
Doshi, G.; Krishnaiyer, R.; Muthukumar, K.;
Parallel Architectures and Compilation Techniques, 2001. Proceedings. 2001 International Conference on , 8-12 Sept. 2001
Pages:257 - 267

[Abstract] [PDF Full-Text (800 KB)] **IEEE CNF**

4 **A concurrent hardware software management scheme for memory hierarchies**
Hammami, O.;
Industrial Electronics, 1994. Symposium Proceedings, ISIE '94., 1994 IEEE International Symposium on , 25-27 May 1994

Pages:368 - 373

[Abstract] [PDF Full-Text (580 KB)] **IEEE CNF**

Home | Log-out | Journals | Conference Proceedings | Standards | Search by Author | Basic Search | Advanced Search | Join IEEE | Web Account | New this week | OPAC Linking Information | Your Feedback | Technical Support | Email Alerting | No Robots Please | Release Notes | IEEE Online Publications | Help | FAQ| Terms | Back to Top

Copyright © 2004 IEEE — All rights reserved

ACM PORTAL
US Patent & Trademark Office
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
memory hierarchy levels
SEARCH
THE ACM DIGITAL LIBRARY
Feedback Report a problem Satisfaction survey

Terms used **memory hierarchy levels** Found **41,045** of **139,988**

Sort results by: publication date
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window
Try an Advanced Search
Try this search in The ACM Guide

Results 121 - 140 of 200 Result page: previous 1 2 3 4 5 6 **7** 8 9 10 next
Best 200 shown Relevance scale

121 An inter-reference gap model for temporal locality in program behavior
Vidyadhar Phalke, Bhaskarpillai Gopinath
May 1995 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1995 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems**, Volume 23 Issue 1
Full text available: pdf(1.16 MB) Additional Information: full citation, abstract, references, citings, index terms

The property of locality in program behavior has been studied and modelled extensively because of its application to memory design, code optimization, multiprogramming etc. We propose a *k* order Markov chain based scheme to model the sequence of time intervals between successive references to the same address in memory during program execution. Each unique address in a program is modelled separately. To validate our model, which we call the Inter-Reference Gap (IRG) model, we show substant ...

Keywords: Markov chains, dynamic memory management, locality of reference, memory replacement, prediction, trace compaction, trace driven simulation

122 Implications of hierarchical N-body methods for multiprocessor architectures
Jaswinder Pal Singh, John L. Hennessy, Anoop Gupta
May 1995 **ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 2
Full text available: pdf(4.66 MB) Additional Information: full citation, abstract, references, citings, index terms, review

To design effective large-scale multiprocessors, designers need to understand the characteristics of the applications that will use the machines. Application characteristics of particular interest include the amount of communication relative to computation, the structure of the communication, and the local cache and memory requirements, as well as how these characteristics scale with larger problems and machines. One important class of applications is based on hierarchical N-body methods, w ...

Keywords: N-body methods, communication abstractions, locality, message passing, parallel applications, parallel computer architecture, scaling, shared address space, shared memory

123 Instruction fetching: coping with code bloat

Richard Uhlig, David Nagle, Trevor Mudge, Stuart Sechrest, Joel Emer
May 1995 **ACM SIGARCH Computer Architecture News , Proceedings of the 22nd annual international symposium on Computer architecture**, Volume 23 Issue 2

Full text available: pdf(1.47 MB) Additional Information: full citation, abstract, references, citings, index terms

Previous research has shown that the SPEC benchmarks achieve low miss ratios in relatively small instruction caches. This paper presents evidence that current software-development practices produce applications that exhibit substantially higher instruction-cache miss ratios than do the SPEC benchmarks. To represent these trends, we have assembled a collection of applications, called the Instruction Benchmark Suite (IBS), that provides a better test of instruction-cache performance. We discuss th ...

124 Energy optimization of multi-level processor cache architectures
Uming Ko, Poras T. Balsara, Ashwini K. Nanda
April 1995 **Proceedings of the 1995 international symposium on Low power design**

Full text available: pdf(86.96 KB) Additional Information: full citation, references, citings, index terms

125 Architectural exploration for datapaths with memory hierarchy
N. D. Holmes, D. D. Gajski
March 1995 **Proceedings of the 1995 European conference on Design and Test**

Full text available: pdf(559.20 KB) Publisher Site Additional Information: full citation, abstract, citings

In this paper, we present a new design-space exploration algorithm, the architecture explorer (AE), for analyzing performance/cost tradeoffs in memory-intensive applications. AE evaluates FU, bus, and memory cost for a series of performance constraints to produce a performance/cost tradeoff curve. Unlike previous approaches, AE handles both hierarchical and non-hierarchical memory architectures with various speeds of memory.

Keywords: architectural exploration, architecture explorer, datapaths, design automation, design-space exploration algorithm, functional units, hardware description languages, high level synthesis, memory architecture, memory architectures, memory cost, memory hierarchy, memory-intensive applications, performance constraints, performance/cost tradeoffs

126 Fine-grain access control for distributed shared memory
Ioannis Schoinas, Babak Falsafi, Alvin R. Lebeck, Steven K. Reinhardt, James R. Larus, David A. Wood
November 1994 **Proceedings of the sixth international conference on Architectural support for programming languages and operating systems**, Volume 29 , 28 Issue 11 , 5

Full text available: pdf(1.20 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper discusses implementations of fine-grain memory access control, which selectively restricts reads and writes to cache-block-sized memory regions. Fine-grain access control forms the basis of efficient cache-coherent shared memory. This paper focuses on low-cost implementations that require little or no additional hardware. These techniques permit efficient implementation of shared memory on a wide range of parallel systems, thereby providing shared-memory codes with a portability ...

127 MOB forms: a class of multilevel block algorithms for dense linear algebra operations
Juan J. Navarro, Toni Juan, Tomás Lang

July 1994 **Proceedings of the 8th international conference on Supercomputing**

Full text available: pdf(1.00 MB) Additional Information: full citation, abstract, references, citings, index terms

Multilevel block algorithms exploit the data locality in linear algebra operations when executed in machines with several levels in the memory hierarchy. It is shown that the family we call Multilevel Orthogonal Block (MOB) algorithms is optimal and easy to design and that using the multilevel approach produces significant performance improvements. The effect of interference in the cache, of the TLB misses, and of page faults are also considered. The multilevel block algorithms are evaluate ...

128 A parallel block implementation of Level-3 BLAS for MIMD vector processors

Michel J. Daydé, Iain S. Duff, Antoine Petitet

June 1994 **ACM Transactions on Mathematical Software (TOMS)**, Volume 20 Issue 2

Full text available: pdf(1.02 MB) Additional Information: full citation, abstract, references, citings, index terms, review

We describe an implementation of Level-3 BLAS (Basic Linear Algebra Subprograms) based on the use of the matrix-matrix multiplication kernel (GEMM). Blocking techniques are used to express the BLAS in terms of operations involving triangular blocks and calls to GEMM. A principal advantage of this approach is that most manufacturers provide at least an efficient serial version of GEMM so that our implementation can capture a significant percentage of the computer performance. A parameter whi ...

Keywords: Level-3 BLAS, matrix-matrix kernels, parallelization, vectorization

129 Tradeoffs in two-level on-chip caching

N. P. Jouppi, S. J. E. Wilton

April 1994 **ACM SIGARCH Computer Architecture News , Proceedings of the 21ST annual international symposium on Computer architecture**, Volume 22 Issue 2

Full text available: pdf(1.16 MB) Additional Information: full citation, abstract, references, citings, index terms

The performance of two-level on-chip caching is investigated for a range of technology and architecture assumptions. The area and access time of each level of cache is modeled in detail. The results indicate that for most workloads, two-level cache configurations (with a set-associative second level) perform marginally better than single-level cache configurations that require the same chip area once the first-level cache sizes are 64KB or larger. Two-level configurations become even more import ...

130 Optimal parallel sorting in multi-level storage

Alok Aggarwal, C. Greg Plaxton

January 1994 **Proceedings of the fifth annual ACM-SIAM symposium on Discrete algorithms**

Full text available: pdf(1.29 MB) Additional Information: full citation, references, citings, index terms

131 A distributed shared memory multiprocessor ASURA: memory and cache architecture

S. Mori, H. Saito, M. Goshima, S. Tomita, M. Yanagihara, T. Tanaka, D. Fraser, K. Joe, H. Nitta

December 1993 **Proceedings of the 1993 ACM/IEEE conference on Supercomputing**

Full text available: pdf(1.17 MB) Additional Information: full citation, references, citings, index terms

132 Concert-efficient runtime support for concurrent object-oriented programming languages on stock hardware
V. Karamcheti, A. Chien
December 1993 **Proceedings of the 1993 ACM/IEEE conference on Supercomputing**
Full text available: pdf(1.01 MB) Additional Information: full citation, references, citings, index terms

133 Predictability of load/store instruction latencies
Santosh G. Abraham, Rabin A. Sugumar, Daniel Windheiser, B. R. Rau, Rajiv Gupta
December 1993 **Proceedings of the 26th annual international symposium on Microarchitecture**
Full text available: pdf(1.51 MB) Additional Information: full citation, references, citings

134 Deterministic distribution sort in shared and distributed memory multiprocessors
Mark H. Nodine, Jeffrey Scott Vitter
August 1993 **Proceedings of the fifth annual ACM symposium on Parallel algorithms and architectures**
Full text available: pdf(1.01 MB) Additional Information: full citation, references, citings, index terms

135 Data locality and load balancing in COOL
Rohit Chandra, Anoop Gupta, John L. Hennessy
July 1993 **ACM SIGPLAN Notices , Proceedings of the fourth ACM SIGPLAN symposium on Principles and practice of parallel programming**, Volume 28 Issue 7
Full text available: pdf(1.36 MB) Additional Information: full citation, abstract, references, citings, index terms

Large-scale shared memory multiprocessors typically support a multilevel memory hierarchy consisting of per-processor caches, a local portion of shared memory, and remote shared memory. On such machines, the performance of parallel programs is often limited by the high latency of remote memory references. In this paper we explore how knowledge of the underlying memory hierarchy can be used to schedule computation and distribute data structures, and thereby improve data locality. Our study i ...

136 Cache write policies and performance
Norman P. Jouppi
May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture**, Volume 21 Issue 2
Full text available: pdf(1.14 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper investigates issues involving writes and caches. First, tradeoffs on writes that miss in the cache are investigated. In particular, whether the missed cache block is fetched on a write miss, whether the missed cache block is allocated in the cache, and whether the cache line is written before hit or miss is known are considered. Depending on the combination of these polices chosen, the entire cache miss rate can vary by a factor of two on some applications. The combination of no- ...

137 On the parallelization of blocked LU factorization algorithms on distributed memory architectures
G. von Laszewski, M. Parashar, A. G. Mohamed, G. C. Fox
December 1992 **Proceedings of the 1992 ACM/IEEE conference on Supercomputing**

Full text available: pdf(906.79 KB) Additional Information: full citation, references, index terms

138 Willow: a scalable shared memory multiprocessor

J. K. Bennett, S. Dwarkadas, J. Greenwood, E. Speight

December 1992 **Proceedings of the 1992 ACM/IEEE conference on Supercomputing**

Full text available: pdf(1.05 MB) Additional Information: full citation, references, citings, index terms

139 Reducing memory latency via non-blocking and prefetching caches

Tien-Fu Chen, Jean-Loup Baer

September 1992 **ACM SIGPLAN Notices , Proceedings of the fifth international conference on Architectural support for programming languages and operating systems**, Volume 27 Issue 9

Full text available: pdf(1.36 MB) Additional Information: full citation, references, citings, index terms

140 Performance evaluation of extended storage architectures for transaction processing

Erhard Rahm

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data**, Volume 21 Issue 2

Full text available: pdf(1.47 MB) Additional Information: full citation, abstract, references, index terms

The use of non-volatile semiconductor memory within an extended storage hierarchy promises significant performance improvements for transaction processing. Although page-addressable semiconductor memories like extended memory, solid-state disks and disk caches are commercially available since several years, no detailed investigation of their use for transaction processing has been performed so far. We present a comprehensive simulation study that compares the performance of these storage ty ...

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player